Bayesian Gait Optimization for Bipedal Locomotion

R. Calandra, N. Gopalan, A. Seyfarth, J. Peters, M. Deisenroth

Lauflabor

Intelligent Autonomous Systems
TU Darmstadt

Locomotion Lab
TU Darmstadt

Outline

- Introduction
- Gait Optimization
- Bayesian Optimization
  - Brief introduction to Gaussian Processes
- Experimental Results
  - LQG
  - Bipedal walker "*Fox*"
- Conclusion

Outline

- Introduction
- Gait Optimization
- Bayesian Optimization
  - Brief introduction to Gaussian Processes
- Experimental Results
  - LQG
  - Bipedal walker “*Fox*”
- Conclusion

Bipedal locomotion

- Versatile kind of Locomotion
- Static vs Dynamic
- Desirable properties:
  - Robust
  - Speed
  - Energy efficient

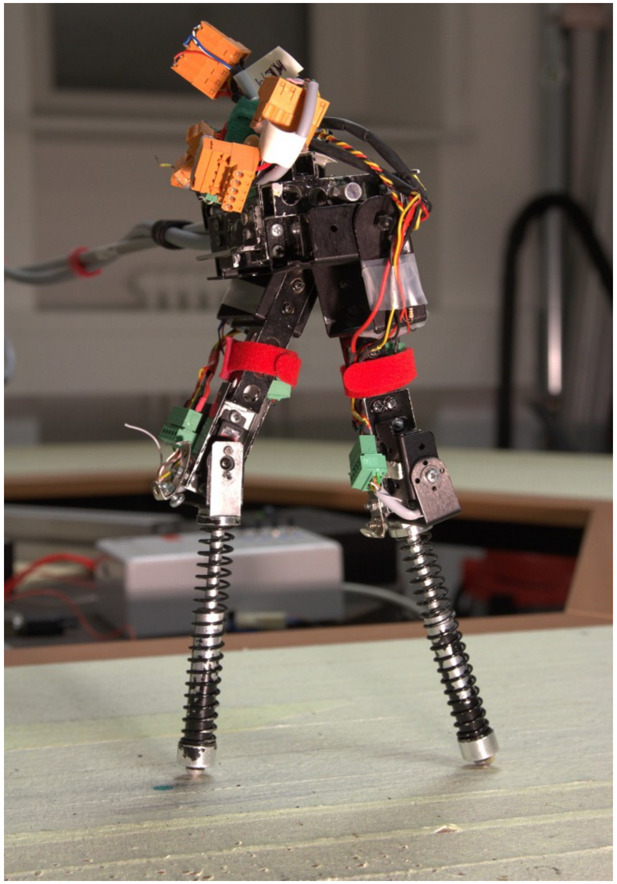

Figure 1: Dynamic walker "Fox"

Gait design

- Controller already exists
- Need to tune parameters:
  - Empirical process
  - Expert-dependent
  - Robot-specific

Gait optimization:
  - Definition of criteria
  - Automatic process

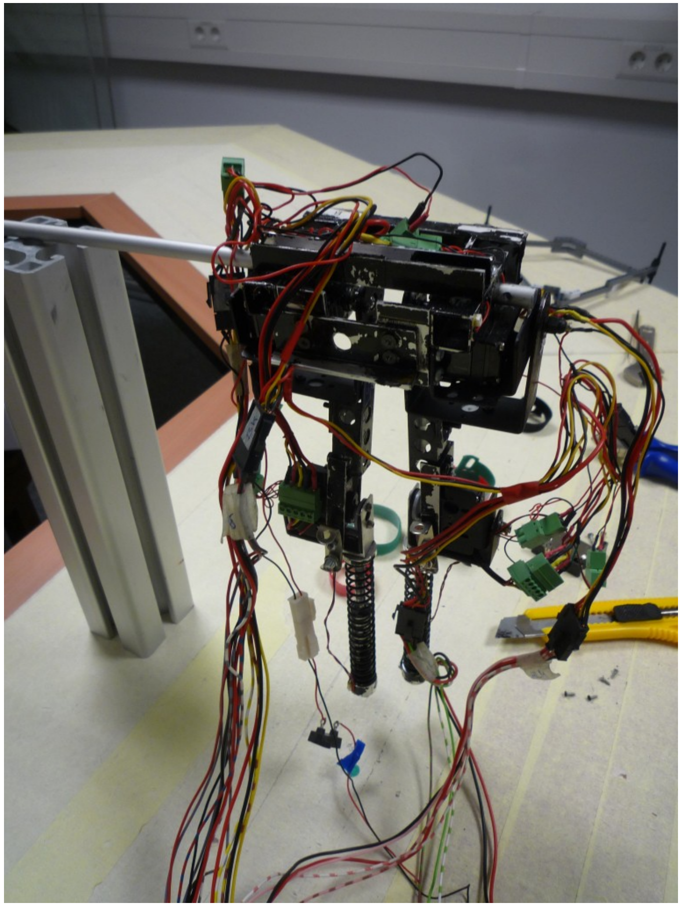

Figure 2: “Fox” disassembled

Outline

- Introduction
- Gait Optimization
- Bayesian Optimization
    - Brief introduction to Gaussian Processes
- Experimental Results
    - LQG
    - Bipedal walker "*Fox*"
- Conclusion

Gait Optimization

$$\text{Minimize}_\theta\ f(\theta)$$

f() = objective function (i.e., criteria to optimize) e.g.:

- *walking speed*
- *energy efficiency*
- *walking robustness*
- *a mixture of the above*

θ *= parameters of the given controller*

Optimization Problem

Necessary properties:

- Global Optimization
- Stochastic objective function (i.e., with noise)
- Zero-order (i.e., no gradients available)

Optimization Problem

Necessary properties:

- Global Optimization
- Stochastic objective function (i.e., with noise)
- Zero-order (i.e., no gradients available)

Desirable properties:

- Limited needs for interactions
- Reusability of past interactions (if available)
- Use of expert knowledge (if available)

Some Optimization Approaches

- Model-free:
  - Grid search
  - Gradient descent
  - Random search
  - Genetic algorithms
  - …

- Model-based:
  - Bayesian optimization
  - ...

PhD student-based optimization

- Student dependent
- Typically requires food

[source: phdcomic.com]

PhD student-based optimization

- Student dependent
- Typically requires food

PhD students prefer to spend time grilling!!!

Grid search

- Create Grid
- Evaluate the points of the grid

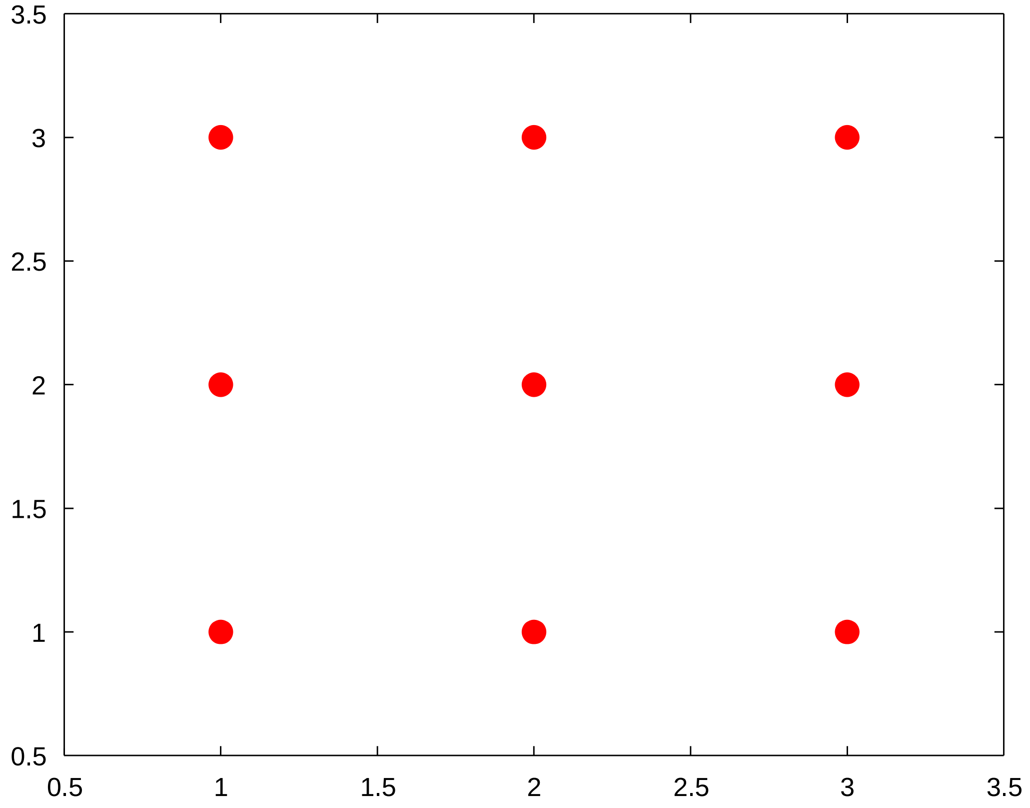

Grid search

- Create Grid
- Evaluate the points of the grid

- Can miss the interesting space!

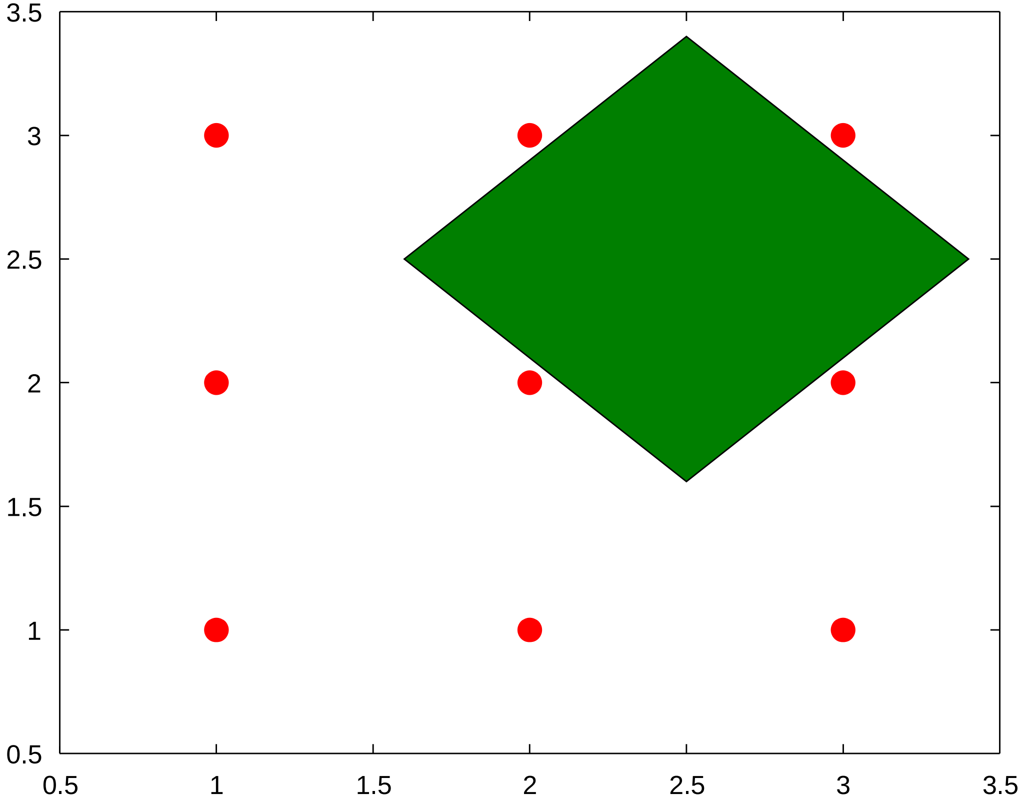

Grid search

N = number of points along each dimension

D = number of parameters

N^D evaluations required !

Grid search

N = number of points along each dimension

D = number of parameters

N^D evaluations required !

Examples:

$N = 2, D = 8 \rightarrow 2^8 = 256$

Grid search

N = number of points along each dimension

D = number of parameters

N^D evaluations required !

Examples:

$N = 2, D = 8 \rightarrow 2^8 = 256$

$N = 3, D = 8 \rightarrow 3^8 = 6561$ → **FORGET ABOUT IT!**

Gradient Descent-based Optimization

$$\theta_{t+1} = \theta_t + \varepsilon \nabla_t$$

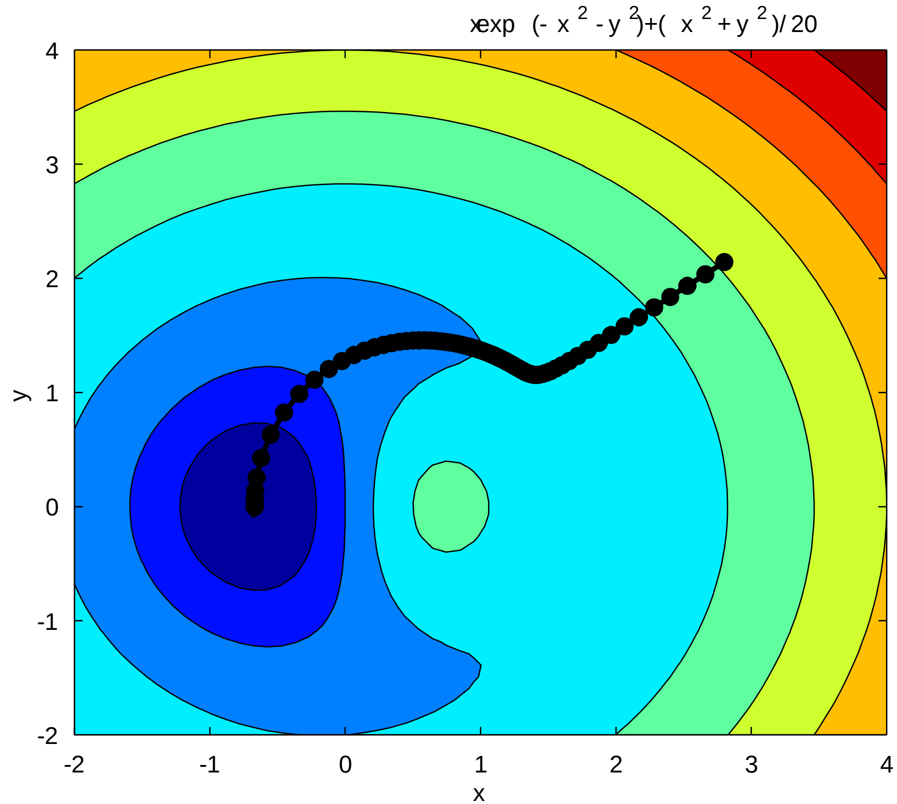

- Local Optimization
- High number of interactions
- No gradients ∇_t available
  - approximate gradient
- No reusability
- No convergence for stochastic

Random Search

$$\theta_{t+1} = \text{rand}()$$

- Global optimization
- Statistical guarantees
- Better than grid search for manifolds [Bergstra 2011]
- High number of interactions required

Model-based Optimization

until *converged*

    Model is created from previous evaluations

    Optimization of the *response surface* (i.e., the model)

    Evaluation of the *proposed* solution

end

Model-based Optimization

until *converged*
    Model is created from previous evaluations
    Optimization of the *response surface* (i.e., the model)
    Evaluation of the *proposed* solution
end

<u>Global optimization "requires" some kind of exploration</u>

Some possible models

- Linear model
- Spline model
- Gaussian Process (GP) model [Kushner 1964]
  - Probabilistic prediction

Outline

- Introduction
- Gait Optimization
- Bayesian Optimization
  - Brief introduction to Gaussian Processes
- Experimental Results
  - LQG
  - Bipedal walker “*Fox*”
- Conclusion

Bayesian Optimization

- Why a probabilistic model?
- How to choose next point to evaluate?
- How to explore?

Regression

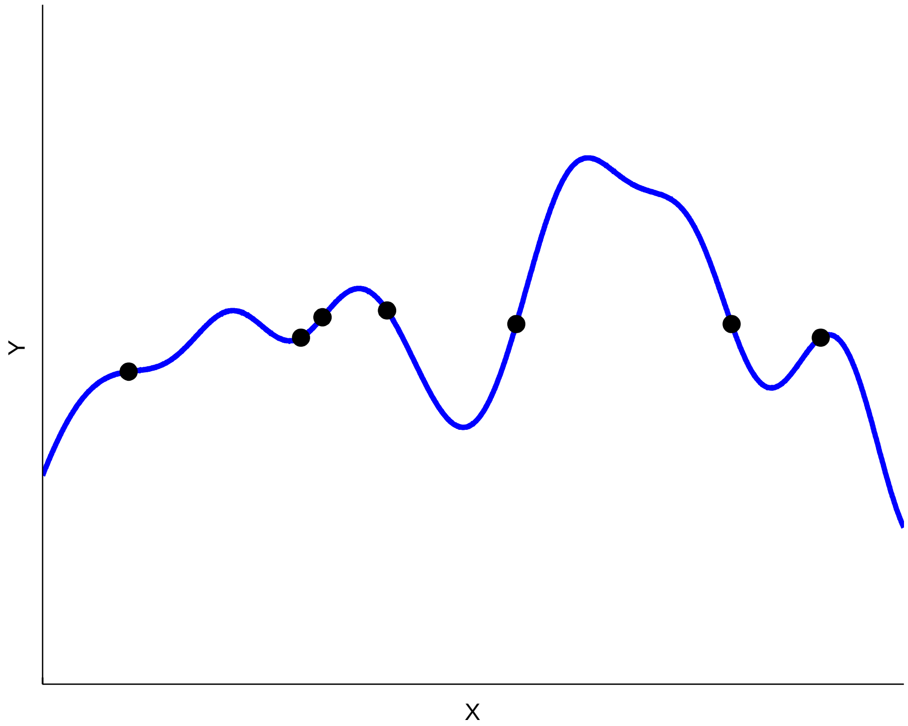

Regression

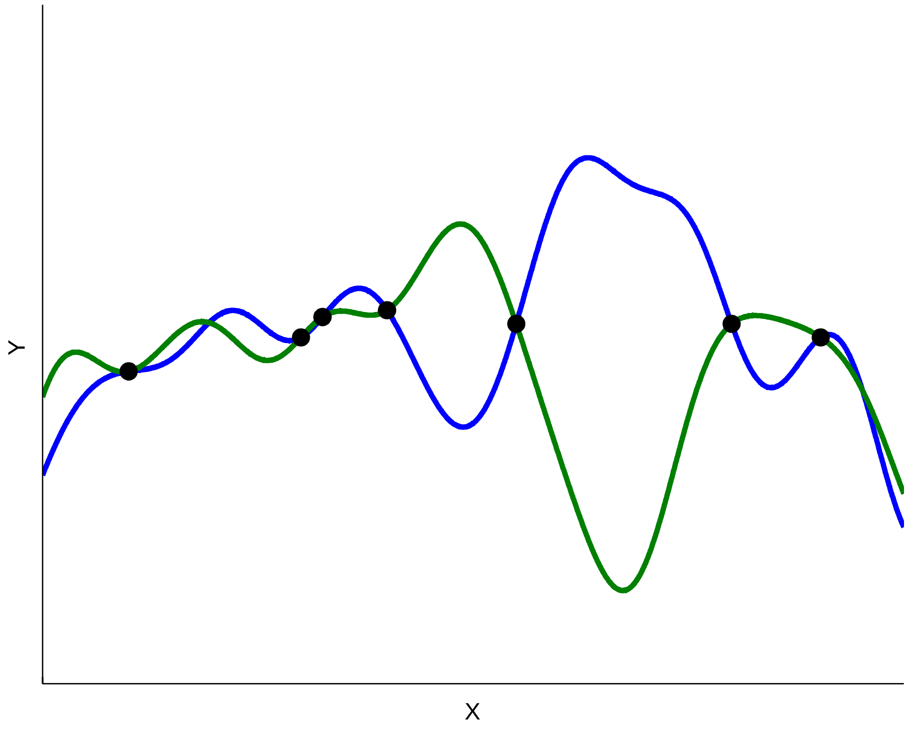

Regression

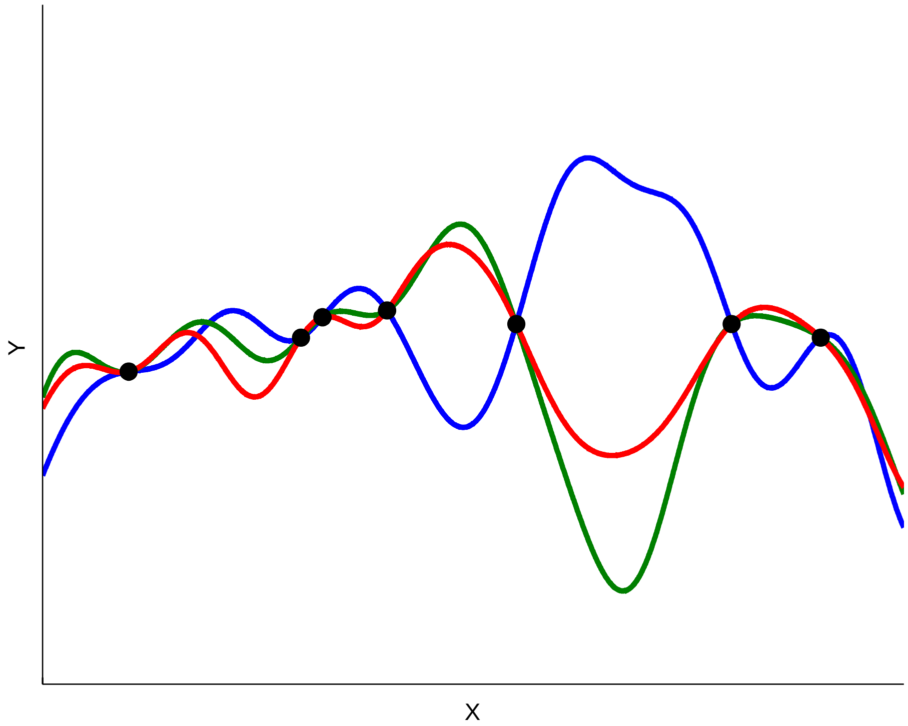

Regression

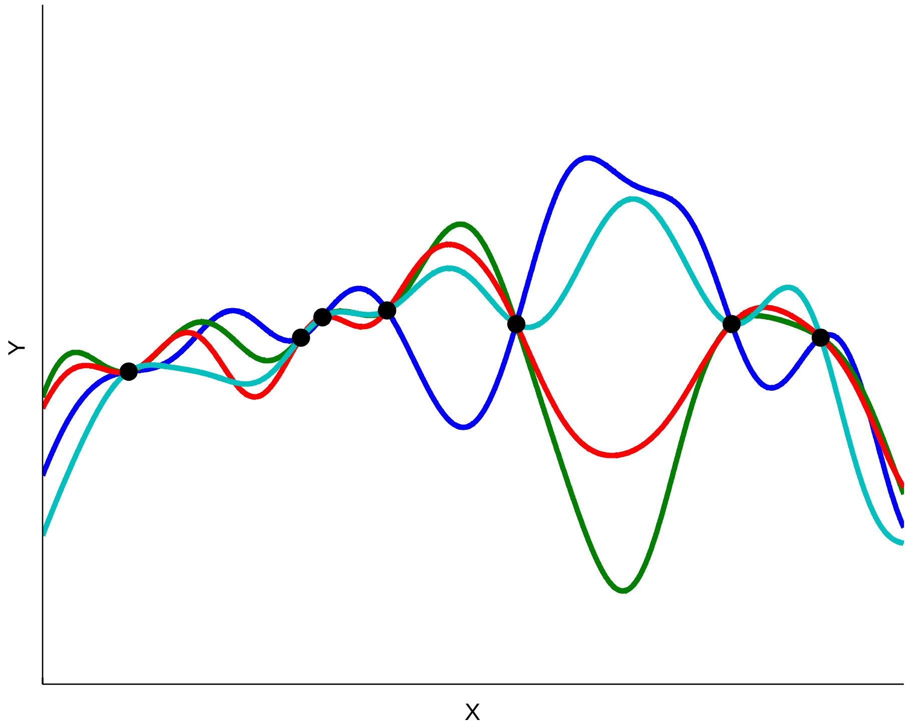

Regression

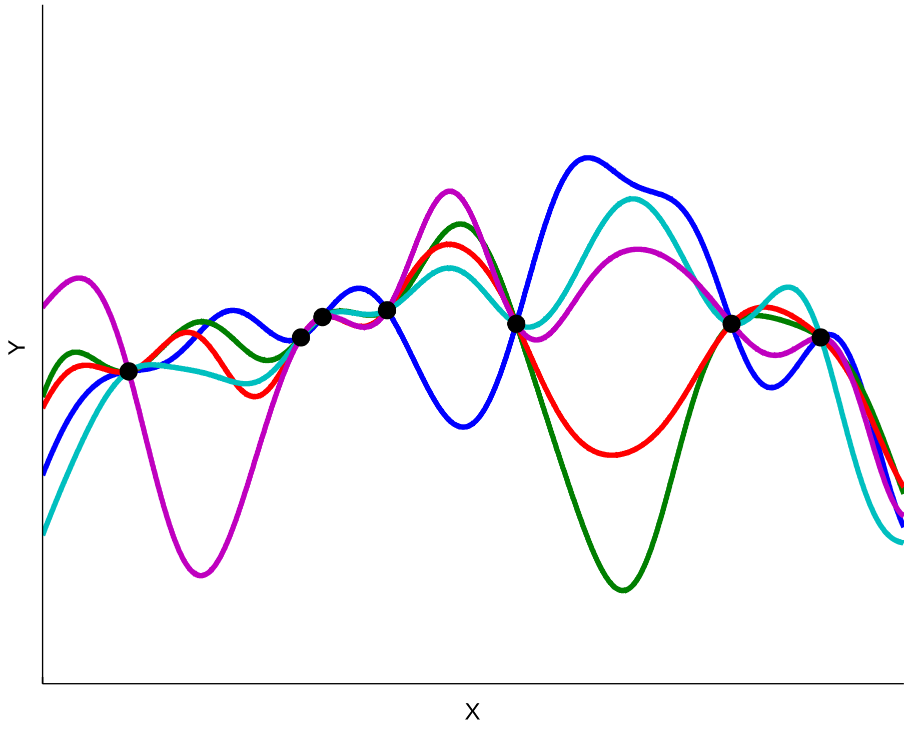

Regression

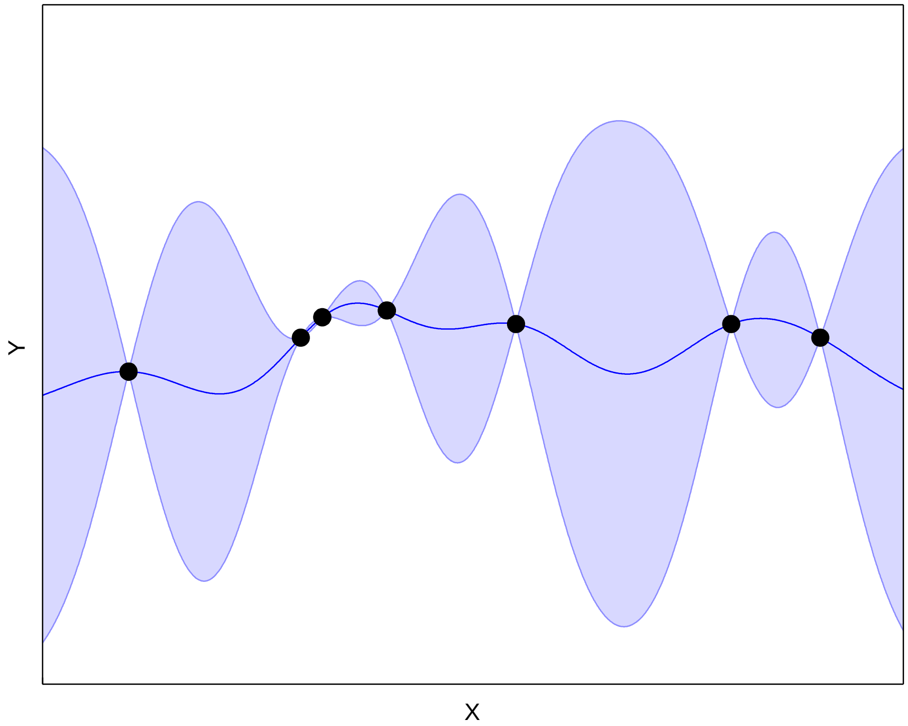

Gaussian Processes modeling

Prediction = probability distribution

- more informative for decision making

Additionally:

- Model is flexible
- Less prone to overfit
- Can model noise

GP Model

$$y = f(\theta) + \epsilon \qquad \epsilon = N(0, \delta^2)$$

$$y(\theta) \begin{cases} \mu(\theta) = k(\theta, X)(K + \delta^2 I)^{-1} y \\ \sigma(\theta) = k(\theta, \theta) - k(\theta, X)(K + \delta^2 I)^{-1} k(X, \theta) \end{cases}$$

Different Covariance functions:

- Square Exponential with Automatic relevance determination (ARD):

$$k(x^p, x^q) = \sigma_f^2 \exp\left(-\frac{1}{2}(x^p - x^q)^T P^{-1}(x^p - x^q)\right)$$

[Rasmussen 2006]

Model Optimization

Acquisition function α: Heuristic criteria that choose the next set of parameters to evaluate

$$\theta_{t+1} = \mathrm{argmin}_{\theta} \; \alpha\left(\mu\left(\theta\right), \sigma\left(\theta\right)\right)$$

- New optimization problem:
  - Global optimization
  - Deterministic
  - First-order: Gradient (and Hessian) in analytical form

Acquisition function

Some acquisition functions α :

- Mean
- Probability of improvement (PI)
- Expected improvement (EI)
- Upper Confidence Bound (UCB) and GP-UCB
- Entropy Search

Trade-offs between exploration and exploitation

Bayesian Optimization

TECHNISCHE UNIVERSITÄT DARMSTADT

```
until converged
    Model is created from previous evaluations
    Computation of the response surface (i.e., the model)
    Optimization of the acquisition surface
    Evaluation of the proposed solution
end
```

<u>Take care of exploration!</u>

Bayesian Optimization Example

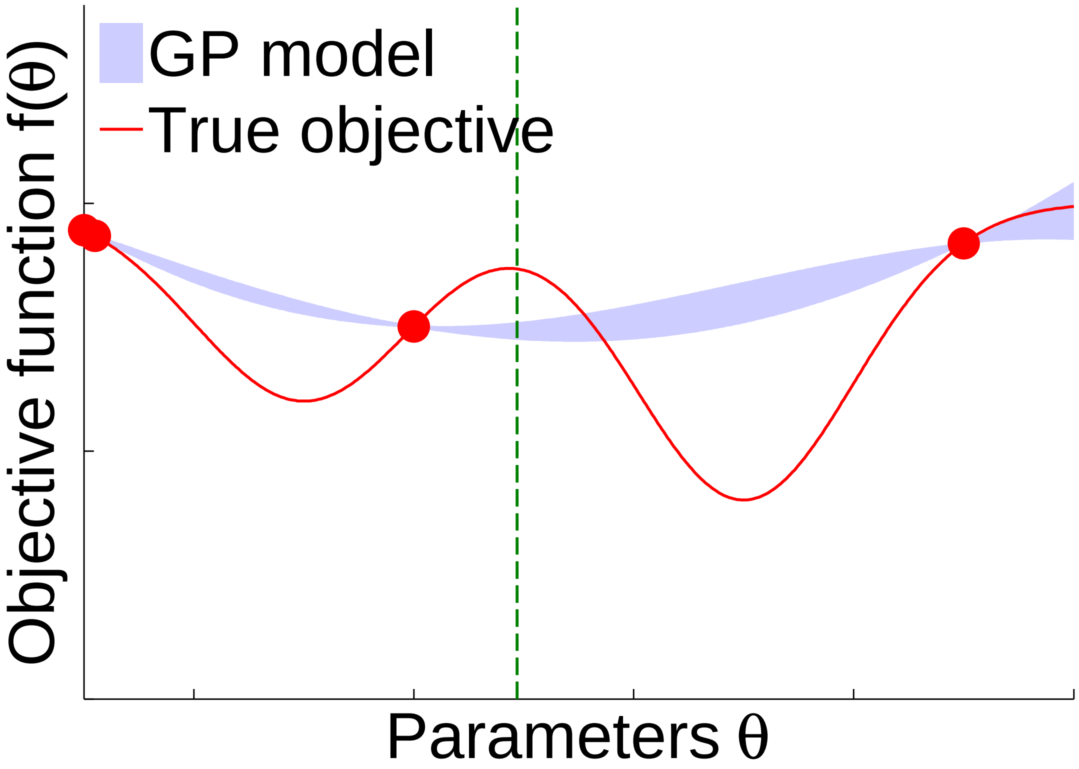

Bayesian Optimization Example

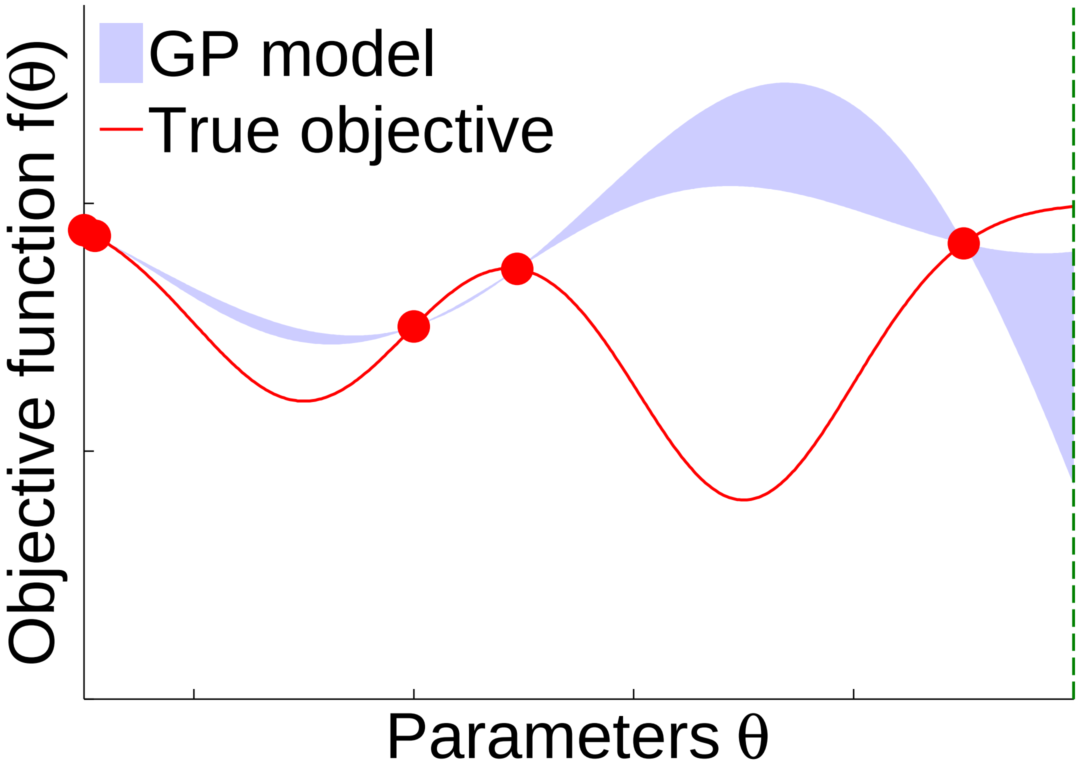

Bayesian Optimization Example

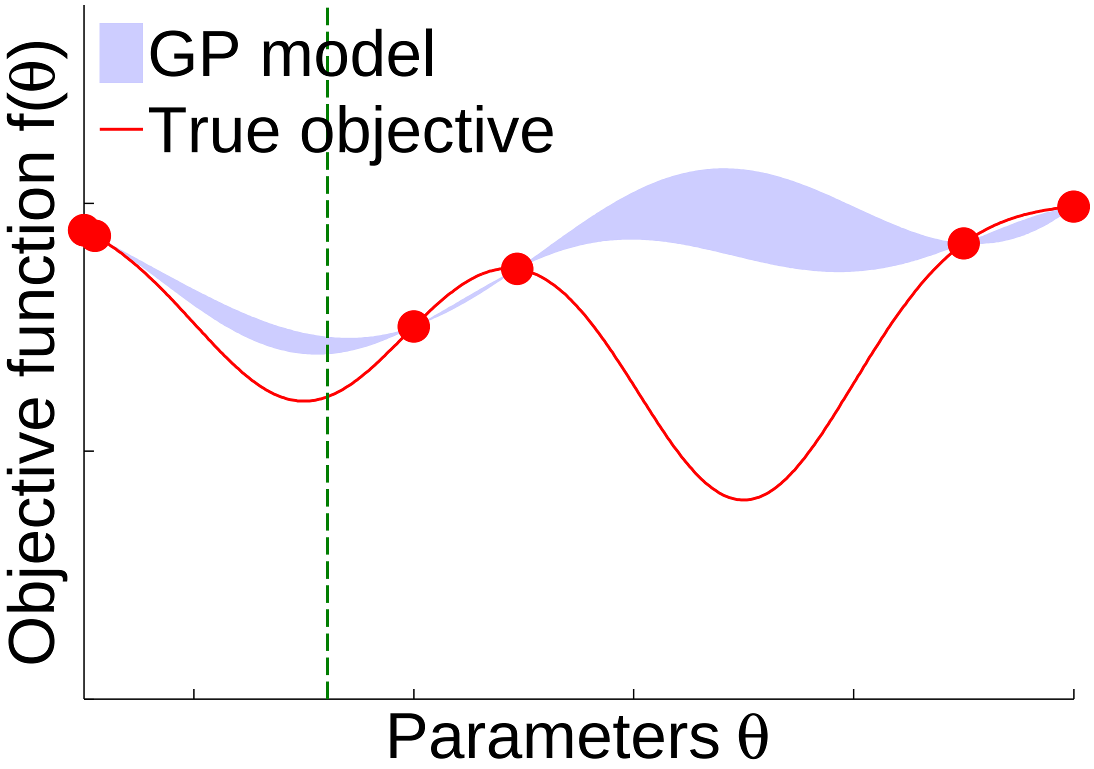

Bayesian Optimization Example

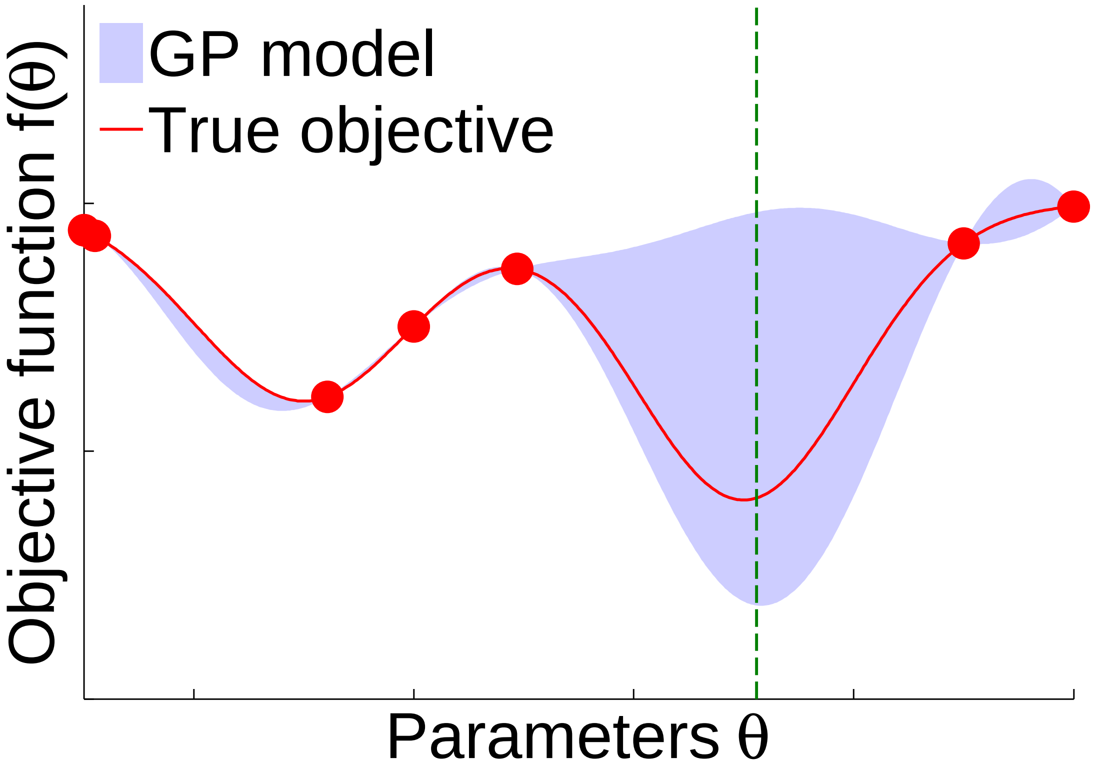

Bayesian Optimization Example

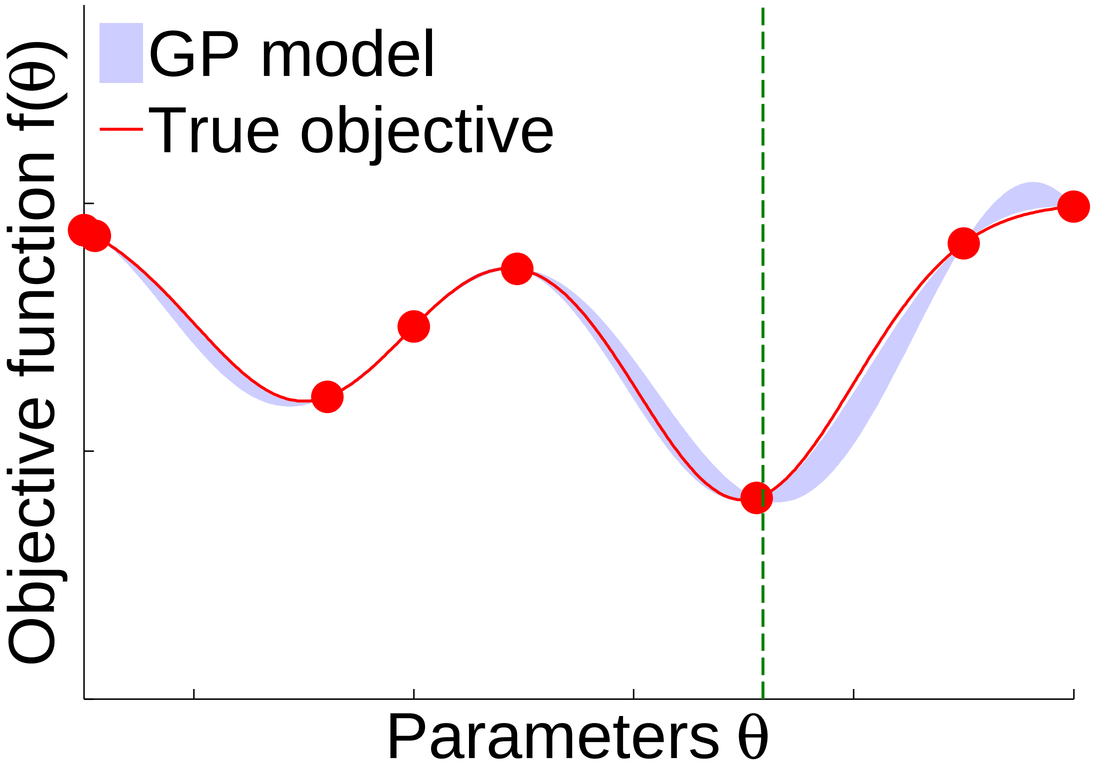

Outline

- Introduction
- Gait Optimization
- Bayesian Optimization
  - Brief introduction to Gaussian Processes
- Experimental Results
  - LQG
  - Bipedal walker “*Fox*”
- Conclusion

Linear-Quadratic-Gaussian Problem

$$x_{K+1} = A x_k + B u_k + w_k \qquad w_k = N(0, \Sigma)$$

$$u_k = L x_k$$

$$J = x'_N Q_N x_N + \sum_{k=0}^{N-1} \left(x'_k Q_k x_k + u'_k R_k u_u \right)$$

- Find gains L that maximize cost J
- (Exists optimal analytical solution)

[Bertsekas 2007]

LQG Optimization

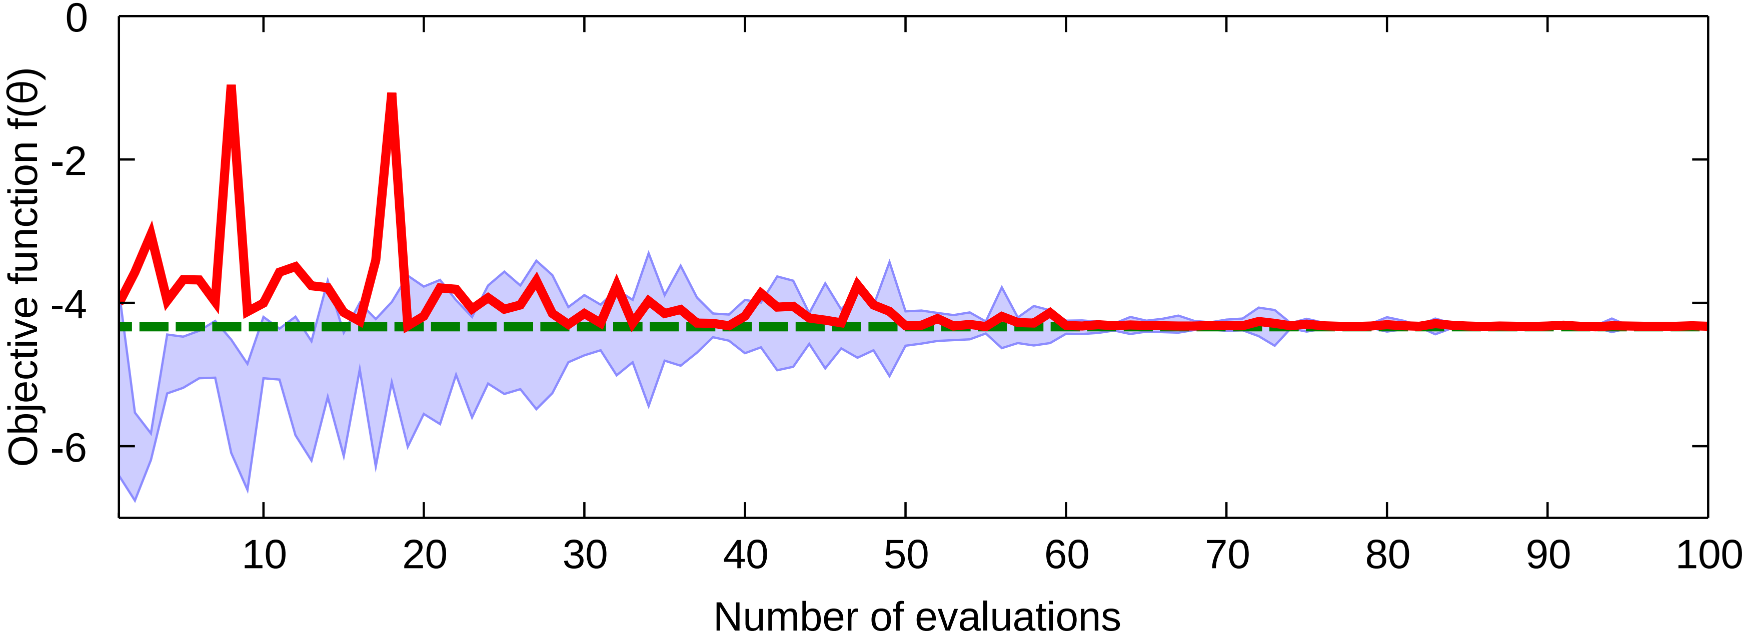

LQG Results

Bayesian Optimization applied to LQG:

- Near-Optimal solutions
- Also applicable to special cases (e.g.,: cases with limited time horizons N, where the algebraic Riccati equation is not applicable, and the discrete time Riccati equation, which can be applied, does not produce a stationary solution)

Fox

- Bipedal Bio-inspired Walker
- 4 actuated DoF:
  - 2 Hips
  - 2 Knees
- Springs in each of the lower legs
- Constrained by a boom to circular walking

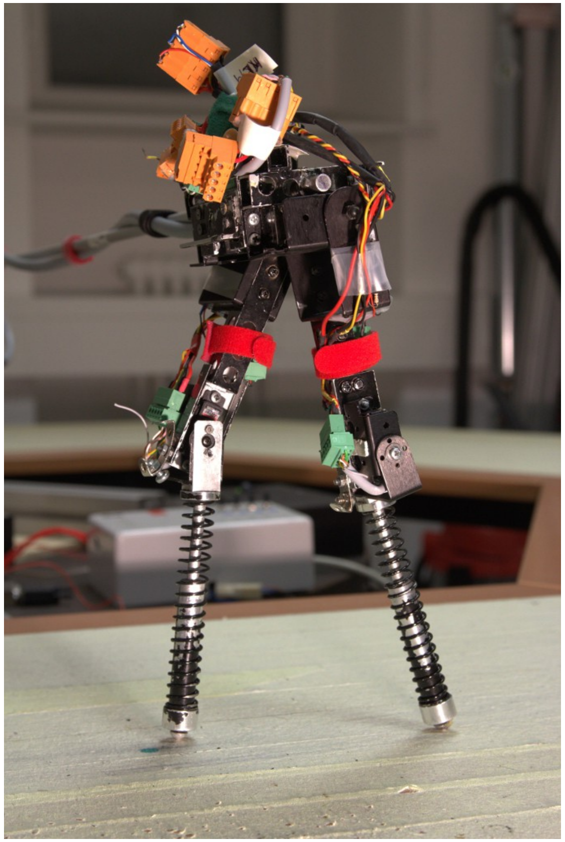

Gait Optimization

Objective function:

$$f(\theta) = -\sum_{i=1}^{3} V_i(\theta)$$

(Intuition: speed + reliability)

Each experiment last 15 sec.

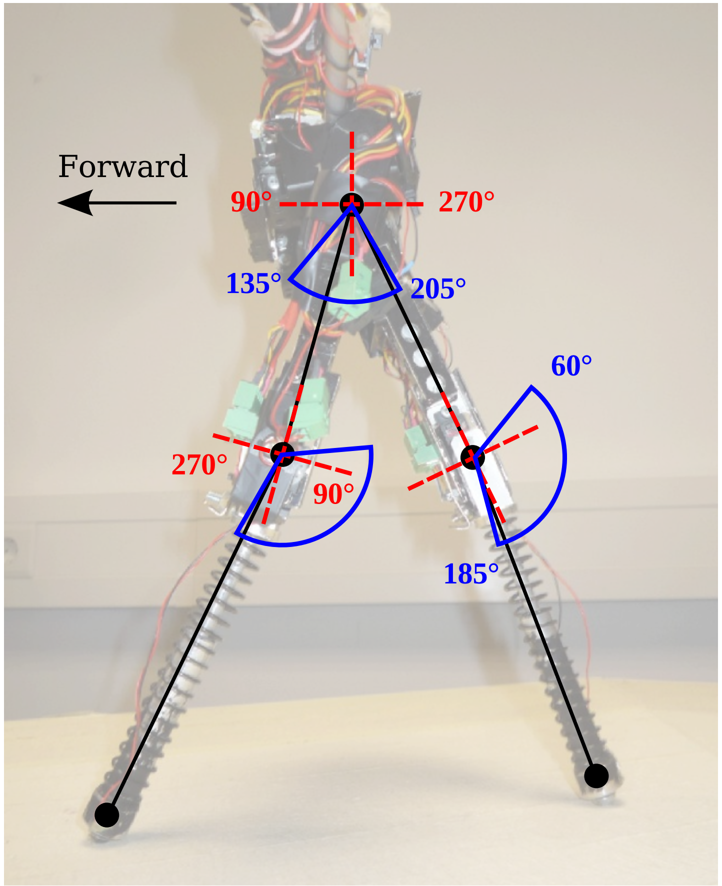

Fox's Controller

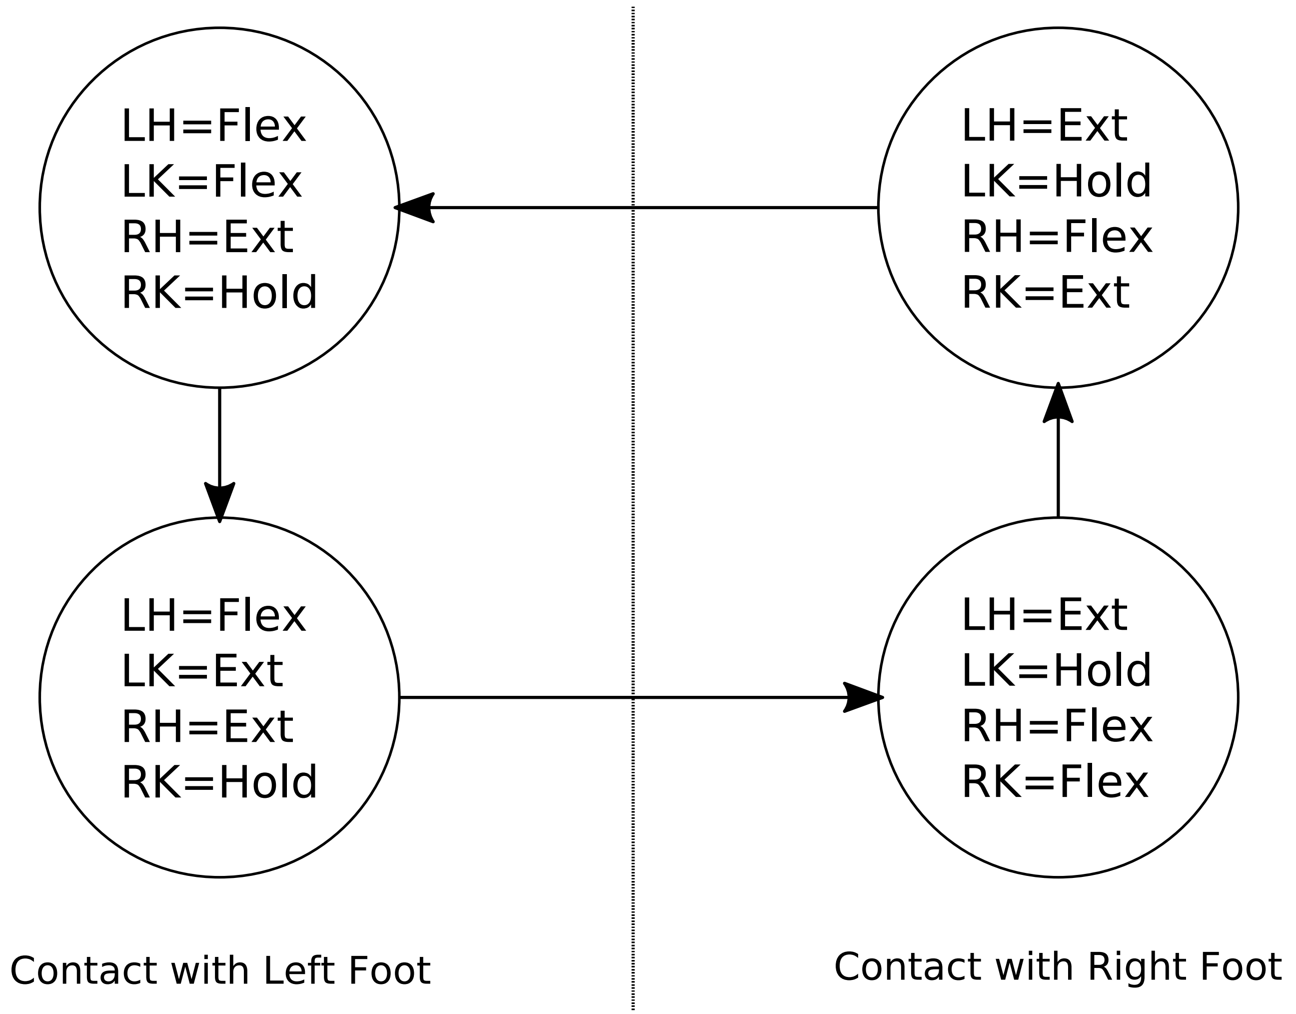

Optimization Parameters

8 Parameters θ:

- 4 transition thresholds between states (i.e., max/min angles of hip joints during walking)
  Asymmetric between legs!
- 4 torques applied during walking:
  - Hip Extension, Hip Flexion
  - Knee Extension, Knee Flexion

Fox's Optimization

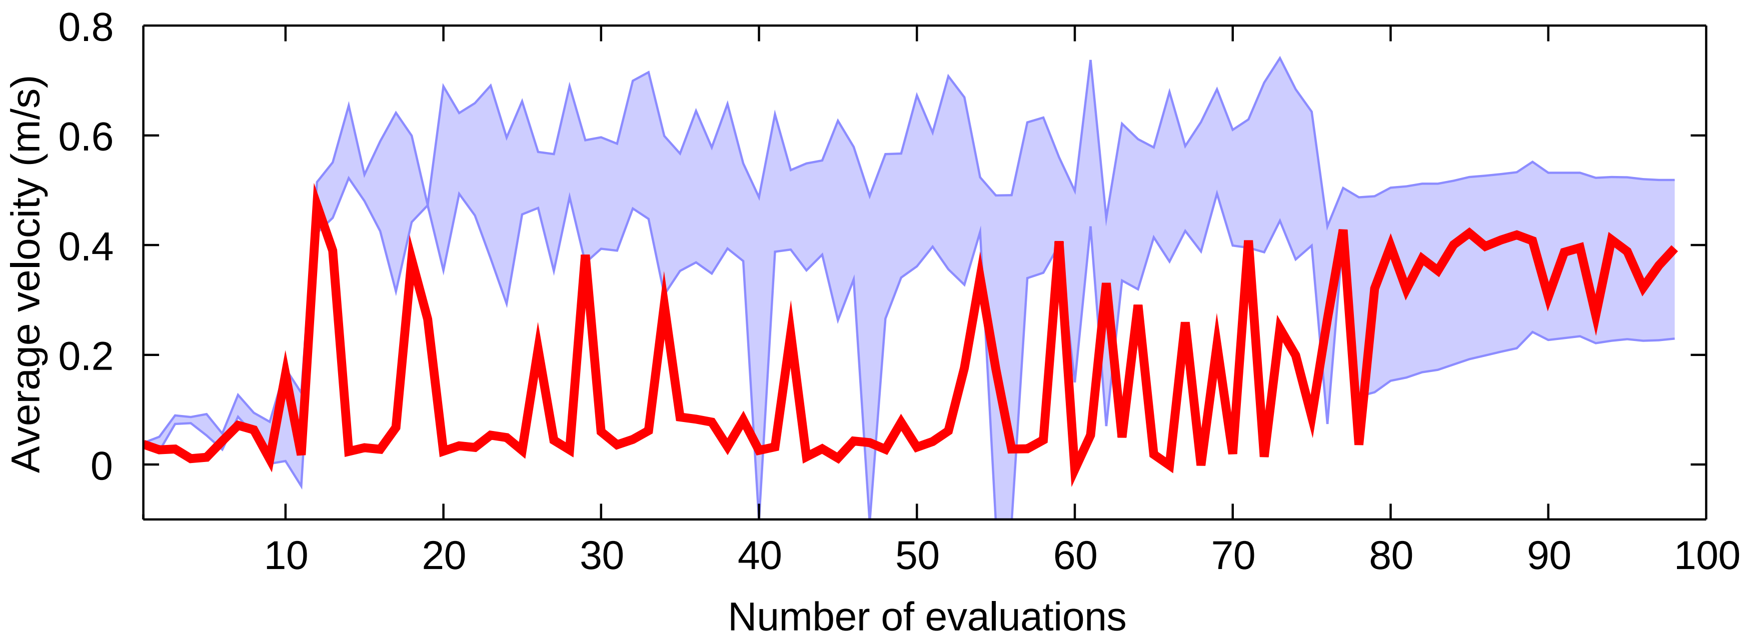

Results of the Optimization

Outline

- Introduction
- Gait Optimization
- Bayesian Optimization
  - Brief introduction to Gaussian Processes
- Experimental Results
  - LQG
  - Bipedal walker "*Fox*"
- Conclusion

Summary

- Gait Optimization
  - Brief overview of optimization methods
- Bayesian Optimization
  - Model-based
  - Global, Stochastic, zero-order optimizer
  - Limited number of interactions required
  - Can use previous experiments and expert knowledge
- Experimental results
  - LQG
  - Bipedal walker FOX

Limitations

Limitations

- No discontinuous value functions
- Boundaries on the parameters to optimize (as in any global optimization problem)
- Low dimensionality of the problem (typically <20)

Difficulties

- Train GP model
- Global optimization

Collaborators

Nakul Gopalan

Marc Deisenroth

Jan Peters

André Seyfarth

So Long, and Thanks for All the Fish

Thank You !

Questions ?

Bibliography

[1] Jones, D. (2001), A taxonomy of global optimization methods based on response surfaces, *Journal of global optimization* **21**(4), 345—383.

[2] Kushner, H. (1964), A new method of locating the maximum point of an arbitrary multipeak curve in the presence of noise, *Journal of Basic Engineering* **86**, 97.

[3] Lizotte, D.; Wang, T.; Bowling, M. & Schuurmans, D. (2007), Automatic gait optimization with Gaussian process regression, *in* 'Proc. of IJCAI', pp. 944—949.

[4] Snoek, J.; Larochelle, H. & Adams, R. (2012), Practical Bayesian Optimization of Machine Learning Algorithms, *arXiv preprint arXiv:1206.2944*.

[5] Tesch, M.; Schneider, J. & Choset, H. (2011), Using response surfaces and expected improvement to optimize snake robot gait parameters, *in* Intelligent Robots and Systems (IROS), 2011 IEEE/RSJ International Conference on', pp. 1069--1074.

Bibliography

[6] Bergstra, J.; Bardenet, R.; Bengio, Y.; Kégl, B. & others (2011), Algorithms for hyper-parameter optimization, *in* '25th Annual Conference on Neural Information Processing Systems (NIPS 2011)'.

[7] Rasmussen, Carl Edward, and Christopher KI Williams. Gaussian processes for machine learning. Vol. 1. Cambridge, MA: MIT press, 2006.

[8] D. P. Bertsekas. Dynamic Programming and Optimal Control. Athena Scientific, 3rd edition, 2007